






Computational Challenges and Directions in the Office of Science

Science for DOE and the Nation
www.science.doe.gov



NCSX






The Office of Science

- Supports basic research that underpins DOE missions.
- Constructs and operates large scientific facilities for the U.S. scientific community.
 - Accelerators, synchrotron light sources, neutron sources, etc.
- Five Offices
 - Basic Energy Sciences
 - Biological and Environmental Research
 - Fusion Energy Sciences
 - High Energy and Nuclear Physics
 - Advanced Scientific Computing Research






Computational Science is Critical to the Office of Science Mission


Scientific problems of strategic importance typically:

- Involve physical scales that range over 5-50 orders of magnitude;
- Couple scientific disciplines, e.g., chemistry and fluid dynamics to understand combustion;
- Must be addressed by teams of mathematicians, computer scientists, and application scientists; and
- Utilize facilities that generate millions of gigabytes of data shared among scientists throughout the world.

The Scale of the Problem



Two layers of Fe-Mn-Co containing 2,176 atoms corresponds to a wafer with dimensions approximately fifty nanometers ($50 \times 10^{-9}\text{m}$) on a side and five nanometers ($5 \times 10^{-9}\text{m}$) thick. A simulation of the properties of this configuration was performed on the IBM SP at NERSC. The simulation lasted for 100 hrs. at a calculation rate of 2.46 Teraflops (one trillion floating point operations per second). To explore material imperfections, the simulation would need to be at least 10 times more compute intensive.




Simulation Capability Needs

FY2004-05 Timeframe




Application	Simulation Need	Sustained Computational Capability Needed (Tflops)	Significance
Climate Science	Calculate chemical balances in atmosphere, including clouds, rivers, and vegetation.	> 50	Provides U.S. policymakers with leadership data to support policy decisions. Properly represent and predict extreme weather conditions in changing climate.
Magnetic Fusion Energy	Optimize balance between self-heating of plasma and heat leakage caused by electromagnetic turbulence.	> 50	Underpins U.S. decisions about future international fusion collaborations. Integrated simulations of burning plasma crucial for quantifying prospects for commercial fusion.
Combustion Science	Understand interactions between combustion and turbulent fluctuations in burning fluid.	> 50	Understand detonation dynamics (e.g. engine knock) in combustion systems. Solve the "soot" problem in diesel engines.
Environmental Molecular Science	Reliably predict chemical and physical properties of radioactive substances.	> 100	Develop innovative technologies to remediate contaminated soils and groundwater.
Astrophysics	Realistically simulate the explosion of a supernova for first time.	>> 100	Measure size and age of Universe and rate of expansion of Universe. Gain insight into inertial fusion processes.




Applications Scientist View






Scientific Discovery Through Advanced Computation (SciDAC)

- SciDAC brings the power of terascale computing and information technologies to several scientific areas -- breakthroughs through simulation.
- SciDAC is building community simulation models through collaborations among application scientists, mathematicians and computer scientists -- research tools for plasma physics, climate prediction, combustion, etc.
- State-of-the-art electronic collaboration tools facilitate the access to these tools by the broader scientific community to bring simulation to a level of parity with theory & observation in the scientific enterprise.





terascale computing

Hardware Infrastructure — Software Infrastructure — Collaboratories and DataGrids




Distribution of Complexity in Scientific Computing





DOE-SC Agency Coordination Overview

	Research Coordination	Development Coordination	Strategy Coordination
NNSA	X – \$17M research funded at NNSA laboratories	X – Red Storm development	X – Formal coordination documents
DOD – DUSD Science and Technology			X – IHEC study
DARPA		X – HPCS review team	X – HPCS evaluation system plan
NSA	X – UPC	X – Cray SV2/X1 development	
All Agencies			X – HECCWG




DARPA High Productivity Computing Systems Program (HPCS)

Goal:





- Provide a new generation of economically viable high productivity computing systems for the national security and industrial user community (2007 – 2010)

Impact:

- **Performance** (efficiency): critical national security applications by a factor of 10X to 40X
- **Productivity** (time-to-solution)
- **Portability** (transparency): insulate research and operational application software from system
- **Robustness** (reliability): apply all known techniques to **protect against outside attacks**, hardware faults, & programming errors



HPCS Program Focus Areas




Applications:

- Intelligence/surveillance, reconnaissance, cryptanalysis, weapons analysis, airborne contaminant modeling and biotechnology

Fill the Critical Technology and Capability Gap

Today (late 80's HPC technology).....to.....Future (Quantum/Bio Computing)



Computing Metric Evolution

Early Computing Metrics

- Clock frequency
- Raw performance (flops)

GHz Race


Current Computing Metrics

- Clock frequency
- Point performance
- Acquisition Price

Tera-flop Race
(Top Ten HPC Centers)

HPCS “Value” Based Metrics

- System performance relative-to-application diversity
- Scalability (flops-to-petaflops)
- Idea-to-solution
- Time-to-solution
- Mean time-to-recovery
- Robustness (includes security)
- Evolvability
- Application life cycle costs
- Acquisition (facilities and equipment) costs
- Ownership (facilities, support staff, training) costs




High End Computing in FY 2004 President's Budget

- Due to its impact on a wide range of federal agency missions ranging from national security and defense to basic science, high end computing—or supercomputing —capability is becoming increasingly critical. **Through the course of 2003, agencies involved in developing or using high end computing will be engaged in planning activities to guide future investments in this area, coordinated through the NSTC.** The activities will include the development of an interagency R&D roadmap for high-end computing core technologies, a federal high-end computing capacity and accessibility improvement plan, and a discussion of issues (along with recommendations where applicable) relating to federal procurement of high-end computing systems. The knowledge gained from this process will be used to guide future investments in this area. Research and software to support high end computing will provide a foundation for future federal R&D by improving the effectiveness of core technologies on which next-generation high-end computing systems will rely.

<http://www.whitehouse.gov/omb/budget/fy2004/pdf/spec.pdf>

Acrobat Page: 181



Issues

- **Many important applications generate poorly structured memory references and are sensitive to latency;**
- **Time to solution begins when the scientist has the idea and ends when the data has been analyzed;**
- **What is the economic model for high(est) performance computers; commercial investment, particle accelerator, or submarine;**
- **How to allocate the complexity and investment between scientific application; mathematical algorithms; systems software; hardware architecture; hardware engineering;....**